

REMARKS

Claims 18, 23-26, 28-29, and 31-33 have been amended and claim 34 has been added. Claims 18, 28, and 31-33 have been amended to address the Examiner's claim objections and rejections under 35 U.S.C. §112, paragraph 2, and 35 U.S.C. §101. Claims 23 and 29 have been amended for clarity and claims 24-26 have been amended to correct antecedent basis as a result of the amendments to claim 18. Claims 19-22, 27, and 30 have not been amended. Upon entry of this amendment, claims 18-34 will be in the application.

The specification has also been amended to add the section headings requested by the Examiner.

Applicant appreciates the Examiner's withdrawal of the Final Rejection dated October 12, 2005, and reopening of prosecution. Applicant also appreciates the Examiner's withdrawal of the objection to the specification and withdrawal of the rejection of the claims over Aihara et al. Reconsideration of the new objections and rejections in view of the above amendments and the following comments is solicited.

Objection to the Disclosure

The Examiner objected to the disclosure for lack of section headings. Applicant has amended the disclosure to add section headings as requested. Withdrawal of the objection to the disclosure is requested.

Objection to Claims 31-32

The Examiner objected to claims 31 and 32 as allegedly being of improper dependent form for allegedly failing to limit claim 18 from which they depend. Claim 31 has been amended to recite a "system having a control apparatus that is programmed to control the objective function of the system according to the method of claim 18," thereby clarifying that claim 31 indeed includes all of the limitations of claim 18. Claim 32 depends from claim 31 and specifies that the claimed system comprises a robot. Withdrawal of the objection to claims 31 and 32 is requested.

Rejection Under 35 U.S.C. §112, 2nd Paragraph

Claims 18-33 stand rejected under 35 U.S.C. 112, second paragraph, as allegedly being indefinite.

The Examiner finds claim 18 “unclear as to how the method recited in the body of the claim accomplishes the controlling of a system, as set forth in the preamble” since the claim “contains no recitation that the candidate action is specifically performed, nor [sic] does it contain any recitation that the performance of a candidate action or choosing a candidate action controls the system per se.” Claim 18 has been amended to recite in a new step d) that the system is commanded “to perform the candidate action identified to be the next performed in step c).” Claim 18, as amended, is believed to overcome the Examiner’s indefiniteness concerns.

The Examiner finds claim 27 “unclear as to how the inclusion of a robot relates to the recited method steps of claim 18.” As noted above, claim 18 has been amended for clarity. Claim 27 recites that the method of claim 18 is performed to control a system where the system comprises a robot. This language is believed to be clear and has not been amended. Claim 27 is believed to be clear and definite as written.

The Examiner finds claim 28 “unclear as to how the method recited in the body of claim 18 would specifically implement and be effected by the limitations of claim 28.” Claim 28 has been amended to reference the ranks of control and how the method controls the system. Claim 28, as amended, is believed to overcome the Examiner’s indefiniteness concerns.

The Examiner finds claim 29 unclear due to its dependency on claim 28. The amendments to claim 28 are believed to overcome the Examiner’s indefiniteness concerns with respect to claim 28. Full examination of claims 28 and 29 in view of the prior art is requested.

The Examiner finds claims 31 and 32 indefinite as “unclear as to what the structure of this system is as claim 18 does not include any system elements” and unclear as to which system the claims refer. Claim 31 has been amended to recite a system “having a control apparatus that is programmed to control the objective function of the system controlled according to the method of claim 18,” thereby overcoming the Examiner’s concerns with respect to claim 31. Claim 32 recites that the system of claim 31 comprises a robot and is thus believed to be definite as well. The amendments to claims 31 and 32 are believed to overcome the Examiner’s indefiniteness concerns with respect to claim 31 and 32.

The Examiner finds claim 33 indefinite as being “unclear as to what the elements of this apparatus are.” Claim 33 has been amended to specifically incorporate the method limitations of claim 18 and to specify that the control apparatus is programmed to perform such method steps. The amendments to claim 33 are believed to overcome the Examiner’s indefiniteness concerns with respect to claim 33.

For the reasons noted above, claims 18-33 as amended and new claim 34 are believed to be clear and definite. If the Examiner still has concerns about the definiteness of the claim language, she is encouraged to contact Applicant’s undersigned representative to discuss claim changes to remove any such ambiguities. Withdrawal of the rejections of claims 18-33 under 35 U.S.C. 112, second paragraph, is solicited.

Rejection Under 35 U.S.C. §101

Claims 18-33 stand rejected under 35 U.S.C. §101 because the claimed subject matter allegedly does not “accomplish a practical application (i.e., a useful, concrete, and tangible result)” in that no “real world outcome” is produced by the claimed method and apparatus since “the chosen action is not used or acted upon.” Applicant disagrees.

Claim 18 as amended recites a method of controlling a system to choose a candidate action to be next performed, and commanding the system to perform the action, effected by repeating specific steps a) to d) to “control the system so as to substantially optimize the objective function of the system.” The tangible result of such a control method is that the resultant objective system functionality may be optimized “to result in the lowest expected growth in regret.” Claim 18 has thus been amended to specify that the chosen action is used or acted upon so as to optimize the objective functionality of the system. Such a method is believed to recite a “practical application” pursuant to the Section 101 Guidelines, and withdrawal of the rejection of claim 18 as not accomplishing a practical application is solicited.

As noted above, claims 31 and 32 have been amended to recite a “system having a control apparatus that is programmed to control the objective function of the system according to the method of claim 18” where the system may comprise a robot. Claims 31 and 32 are believed to produce a practical application for the same reasons as noted above with respect to claim 18.

In view of the above, claims 18-33 and new claim 34 are believed to meet the statutory requirements patentability under 35 U.S.C. §101 and to satisfy the Section 101 Guidelines. If the Examiner still has concerns about whether any of claims 18-34 recites a “practical application” as worded, she is encouraged to contact Applicant’s undersigned representative to discuss claim changes to provide the requisite “practical application.” Withdrawal of the rejection of claims 18-33 under 35 U.S.C. §101 is solicited.

Rejections Under 35 U.S.C. §103

In the Official Action, claims 18-21, 24-25, 30-31, and 33 were rejected under 35 U.S.C. §103(a) as allegedly being obvious over Merriman et al. (US 2002/0099600) in view of Eppen et al. (“Quantitative Concepts for Management”). In addition, claims 22 and 26 were rejected under 35 U.S.C. §103(a) as allegedly being obvious over Merriman et al. and Eppen et al. further in view of McClave et al. (“A First Course in Business Statistics”); claim 23 was rejected under 35 U.S.C. §103(a) as allegedly being obvious over Merriman et al. and Eppen et al. further in view of Jameson (US 6,032,123); and claims 27 and 32 were rejected under 35 U.S.C. §103(a) as allegedly being obvious over Merriman et al. and Eppen et al. further in view of Strickland et al. (US 5,790,407). These rejections are believed to be improper and are respectfully traversed.

As noted in the previous amendment response, the claimed method and system implement the concept of “regret” where “regret is a term that represents a system performance measure that considers the relative merit of exploration of one or more apparently non-best candidate actions, with respect to the relative merit of exploiting what appears to be the current best candidate action based on historical response performances to date.” A system is controlled to optimize its objective function by choosing among a plurality of candidate actions to minimize the growth of regret as so defined, performing the chosen candidate action, and monitoring response performances of a performance of the chosen candidate action. For example, the method of claim 18 includes the steps of:

- a) monitoring response performance of a respective candidate action that is chosen to be performed by the system;
- b) storing, according to candidate action performed by the system, a representation of said monitored response performance;

c) choosing which of the plurality of candidate actions is next performed by the system so as to optimize said objective function by assessing, using the probability distribution of the response performance of all of said plurality of candidate actions, which candidate action is estimated to result in the lowest expected growth in regret after the chosen candidate action is performed by the system;

d) commanding the system to perform the candidate action identified to be the next performed in step c); and

e) repeating steps a) to d) to control the system so as to substantially optimize the objective function of the system.

Such claimed features are not believed to be taught or contemplated by any of the references cited by the Examiner.

The prior art cited by the Examiner falls into two main categories:

1. Systems that contain a controller which makes decisions, stores the results, then makes improved decisions through some analysis of historical decisions and the corresponding outcomes; and

2. Texts relating to methods which recite the term "regret."

Applicant submits that all of the cited prior art systems suffer from a set of common problems that prevent them from being able to operate efficiently, and in a truly automated and self-regulating way under conditions which are not stable over time. The claimed system and method overcomes the limitations of prior art systems and is designed to operate with optimum efficiency in demanding real-world applications, such as those where:

1. The relationships between actions and outcomes change over time (a tool may become worn in a production application, or market conditions may change in a targeted advertising application say). This creates a need for continuous and efficient ongoing learning of observed relationships.

2. The conditions under which one makes each decision are characterized by a large number of variables, but one has no prior knowledge about how those various descriptors relate to the likely success or failure of a particular action under any given set of conditions. As such, both the learning of the relationships between actions and outcomes

under particular circumstances, and the exploitation of any learned knowledge, must happen simultaneously and continuously for the system to operate efficiently.

3. Because of the presence of randomness in the nature of response outcomes to actions taken, those outcomes that have been observed from particular actions in the past may not be typical and may therefore be unrepresentative of the true performance of these actions under those given conditions.

4. The available set of candidate actions (from which the system may select) may change at any time.

The problems associated with prior art methods are explained in the specification from paragraphs [0044] to [0059] inclusive.

None of the prior art documents cited by the Examiner discusses the requirements of learning efficiency as the cited prior art references do not attempt to address nor recognize the role of learning efficiency in the overall performance of self-regulating decision systems. The claimed method and system provides a true self-regulating system that operates with the highest possible learning efficiency (a characteristic which is expressed specifically in the claims as resulting in the "lowest expected growth of regret") given there exists randomness in the nature of the response outcomes. It can do so in the face of changing conditions, even where the interaction scenarios are characterized by a very large number of descriptors. By doing so, the claimed system and method overcomes many of the problems of existing systems that compromise or limit their useful real-world applicability.

These distinctions are particularly apparent when examining the references applied against the claims.

As noted by the Examiner, Merriman et al. disclose a system and method of controlling a system to optimize an objective function by selecting a candidate action, monitoring the system performance in response to the selected candidate action, storing a representation of the monitored response performance, choosing the next candidate action to optimize the objective function, and repeating the process. Thus, the Merriman et al. system makes decisions, stores the results and attempts to make improved decisions based on historical observations in a manner consistent with the prior art systems noted in the background portion of the specification. The Merriman et al. system thus has the same limitations as these prior art systems. Moreover, as specifically acknowledged by the

Examiner (page 9, lines 5-11, of Official Action), Merriman et al. do not teach that the objective function is optimized by assessing the probability distributions of all the candidate actions in order to control the growth of regret, “where regret is a term that represents a system performance measure that considers the relative merit of exploration of one or more apparently non-best candidate actions with respect to the merit of exploring what appears to be the current best candidate action.” For such a teaching, the Examiner refers to Eppen et al. However, Eppen et al. also fail to provide such a teaching, so the obviousness rejection must fail for lack of *prima facie* obviousness.

Eppen et al. explains his concept of “regret” on page 511 of the cited text as “synonymous with the ‘opportunity cost’ of not making the best decision given a state of nature.” Eppen et al. continue to explain that in the example given that “If [the decision maker] knows a probability distribution on the state of nature, he could minimize the expected regret.” Those skilled in the art will appreciate that none of the examples provided by Eppen et al. relate to conditions for minimizing the growth of “regret” under which the claimed system and method operates. In particular, the examples provided by Eppen et al. do not consider the merit of actually taking an action which may be expected to offer an apparently lower immediate payoff because the value of the new information that would be gained by such an action may exceed the loss resulting from deliberately selecting an action which appeared to offer a lower pay-off. This is known as an exploration-exploitation trade-off and is specifically incorporated into the claimed concept of “regret.” Such a notion of “regret” is not taught by Eppen et al.

In all of the examples provided by Eppen et al., it is assumed that (probabilistic) information about the state of nature or the interaction environment is obtained or obtainable by means other than the decision process itself (and its consequent actions and observations). Thus, Eppen et al. provide a standard model of decision making under uncertainty, where the expected payoff criterion is appropriate.

In contrast, the claimed system and method relates to a scenario in which there is no assumption that external information is available, and in which the necessary information is described as being collected in the form of sets of observations about responses and response conditions. In the general environment of the present application, both the learning of the relationships between actions and outcomes under particular circumstances, and the

exploitation of any learned knowledge generally happen simultaneously and continuously for the system to operate efficiently.

The problem addressed by the claimed notion of “regret” is thus one in which one must accept that while obtaining information by making certain decisions, one might lose payoffs achievable by making other decisions. To implement such a notion of “regret” requires a special decision control to make decisions that trade off the benefits of additional information with the losses of ignoring apparently higher immediate payoffs. This is made even more complex in many real-world situations where the interaction scenario may be characterized by a large number of descriptor variables. In such a case, one must operate this optimal trade-off between exploration and exploitation while simultaneously learning the multivariate relationships between actions and expected outcomes under any given set of conditions. These concepts are not expressed by Eppen et al., nor are any solutions to this type of problem discussed.

Eppen et al. summarizes on page 520 (Section 14.6 “A Mid-Chapter Summary”) the three types of conditions for which he describes solutions. In this section, Eppen et al. describe the three cases studies as:

- Decisions under certainty. The decision maker knows exactly what state of nature will occur. His ‘only’ problem is then to select the best decision. Eppen et al. describes this as a deterministic problem under these conditions.
- Decisions under risk. Where a probability distribution is specified on the states of nature.
- Decisions under uncertainty. Where it is assumed that the decision maker has no knowledge about which state of nature will occur.

In all of these three types of problem, Eppen et al. assumed that the payoff is a known quantity. In the situations described by the claimed notion of “regret,” no such assumption is made. In fact, the only way of knowing what potential payoff might be available would be to test each action under every candidate state of nature. It must also be noted that an iterative application of any of Eppen et al.’s three solutions does not result in an exploration-exploitation trade-off of the type required to control the growth of regret as claimed. Accordingly, Applicant submits that Eppen et al. do not address the type of problem to which

the present invention is directed and clearly do not disclose minimizing the growth in “regret” as claimed, where such “regret” is defined as:

a term that represents a system performance measure that considers the relative merit of exploration of one or more apparently non-best candidate actions, with respect to the relative merit of exploiting what appears to be the current best candidate action based on historical response performances to date.

As set forth in M.P.E.P. §§2142-2143.03, in order to establish a *prima facie* case of obviousness, patent examiners are required to establish three criteria: (1) there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings; (2) there must be a reasonable expectation of success; and (3) the prior art reference, or combination of references, must teach or suggest all the claim limitations. The examiner bears the initial burden of factually supporting any *prima facie* conclusion of obviousness. To make a proper obviousness determination, the examiner must “step backward in time and into the shoes worn by the hypothetical ‘person of ordinary skill in the art’ when the invention was unknown and just before it was made.” In view of the available factual information, the examiner must make a determination as to whether the claimed invention “as a whole” would have been obvious at that time to a person of ordinary skill in the art. Importantly, a rejection based on these criteria must be based on what is taught in the prior art, not the applicant’s disclosure. The applicant’s disclosure may not be used as a blueprint from which to construct an obviousness rejection.

In view of the fact that Merriman et al. and Eppen et al. taken together do not teach the claimed minimization in the growth of “regret,” and hence do not teach all of the claimed features of the invention, even if the teachings of Merriman et al. could be combined with the teachings of Eppen et al. as the Examiner suggests, the invention of independent claim 18 would not result. The Examiner has thus failed to establish *prima facie* obviousness and the rejection of claim 18 and all claims dependent thereon should be withdrawn.

Moreover, the Examiner has further failed to provide a *prima facie* case of obviousness with respect to any claim since the Examiner has not met her burden of

providing a suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to combine the reference teachings. Instead, the Examiner has mistakenly stated that “it would have been obvious to one of ordinary skill in the art at the time of the invention to use probability distributions and the theory of regret in the iterative predictive model of Merriman et al. in order to increase the efficiency of utilizing advertising/action space by providing a decision framework with which to analyze the various options.” The general citations provided by the Examiner do not provide the requisite teachings, suggestions or motivations to combine the teachings of Merriman et al. and Eppen et al. in the manner contemplated by the Examiner. In any case, the proposed combination clearly does not suggest the claimed system and method for minimizing the growth of regret in controlling a system. As a result, one skilled in the art would not be motivated to combine the teachings of Merriman et al. and Eppen et al. to provide a system and method for minimizing the growth of regret as claimed.

In view of the above, if the Examiner elects to maintain the obviousness rejections of independent claim 18, the Examiner is strongly urged to clearly articulate the evidence of suggestions, motivations, or knowledge possessed by those skilled in the art that would have led one skilled in the art to combine the teachings of the cited references to arrive at the claimed invention. In the absence of the requisite teachings and motivations to combine teachings to establish *prima facie* obviousness, the rejections of claims 18-33 as being obvious over Merriman et al. and Eppen et al. or any other prior art reference is improper and withdrawal of the obviousness rejections is respectfully solicited.

The Examiner introduces McClave et al. with respect to claims 26 and 32 as an example of the use of the student's t parameter and associated test as a measure of uncertainty of risk in appraising performance given a limited sample of observations. However, McClave et al. do not discuss the use of the student's t parameter as a means to regulate the exploration and exploitation trade-off which is a requirement to deliver the lowest expected growth of regret as claimed. McClave et al. do not discuss nor solve for conditions where investment in learning about outcomes and learning the relationships between outcomes, actions and states of nature is itself part of the problem. The complexities involved in the generic management of the growth of regret, as claimed, are that no prior information is assumed, and that both data from which learning must occur, as well as any potential payoffs,

are dependent upon which actions are taken. Therefore, to learn the (often complex multivariate) relationships between actions, outcomes and interaction scenario descriptors, one requires a special decision control to make decisions that trade off the benefits of gaining additional information with the losses of ignoring apparently higher immediate payoffs under any given set of prevailing conditions. It is this decision process that enacts the claimed control of the growth of regret. While the student's t parameter under certain circumstances may be used as a useful characterizer of risk, its use as a parameter within a complex system to minimize the growth of regret as claimed is not discussed. Thus, while McClave et al. introduce the concept of the student's t parameter as a measure of uncertainty, McClave et al. do not discuss solutions relevant to advance the development of a system which explicitly controls the expected growth of regret as claimed in Claim 18 of the present application. Accordingly, even if one skilled in the art would have been motivated to combine the teachings of McClave et al. with the teachings of Merriman et al. and Eppen et al. as the Examiner asserts, the claimed invention could not have resulted.

The Examiner cites Jameson with respect to claim 23 as teaching maximizing a return given resource constraints by assuming that outcome values can be predicted precisely given a set of conditions that define costs, potential demand, unallocated resource opportunity costs and other factors. The Jameson system has access to information which does not have to be gained by taking specific actions. The Monte Carlo process taught by Jameson is dependent upon information that is available before the process starts, from external sources. At no stage in Jameson's process does the claimed system make decisions that trade-off the value of potentially acquiring new information with the potential losses that are realized by ignoring other candidate actions in a way which results in the lowest expected growth of regret. Accordingly, even if one skilled in the art would have been motivated to combine the teachings of Jameson with the teachings of Merriman et al. and Eppen et al. as the Examiner asserts, the claimed invention could not have resulted.

The Examiner cites Strickland et al. with respect to claims 27 and 32 as teaching a control system for controlling robots. However, the control mechanism disclosed by Strickland et al. does not take actions that result in the lowest expected growth of regret as claimed in claim 18. The environment in which Strickland's disclosed system operates is one in which outputs are measured with respect to predefined response profiles of an external

device (such as a robot) for the purposes of detecting and controlling failures. This environment is not one in which information about outcomes can only be obtained by taking specific discrete actions. Strickland et al. have no need to regulate an exploration and exploitation trade-off in the cited decision process, and no such trade-off condition is identified as existing (nor does one appear to exist in the Strickland et al. disclosure). Using regulation of the exploration and exploitation trade-off is a requirement to effect explicit control in the growth of regret as claimed in claim 18. In the absence of such teachings by Strickland et al., the teachings of Strickland et al. are not believed to be particularly relevant to the claimed system and method.

Accordingly, none of the prior art references cited by the Examiner teaches the concept of “regret” for controlling the objective functioning of a system as claimed, where “regret is a term that represents a system performance measure that considers the relative merit of exploration of one or more apparently non-best candidate actions, with respect to the relative merit of exploiting what appears to be the current best candidate action based on historical response performances to date.” On the contrary, prior art systems are generally limited to making decisions using known historical data without trading off exploring new actions that may lead to better results versus exploiting the apparent current best candidate action. This ability to optimally control the exploration for better solutions is set forth in all independent claims and clearly distinguishes the claimed invention from the prior art control systems cited by the Examiner. Withdrawal of all prior art rejections is solicited.

Conclusion:

For the above reasons, it is submitted that the present application is in condition for allowance and a Notice of Allowability is solicited.

DOCKET NO.: ASTB-0044
Application No.: 09/814,308
Office Action Dated: April 5, 2006

Date: July 21, 2006

PATENT

A handwritten signature in cursive script, reading "Michael P. Dunnam", written over a horizontal line.

Michael P. Dunnam
Registration No. 32,611

Woodcock Washburn LLP
One Liberty Place - 46th Floor
Philadelphia PA 19103
Telephone: (215) 568-3100
Facsimile: (215) 568-3439